arXiv:1409.3742v1 [cs.DS] 12 Sep 2014

Data Reductions and Combinatorial Bounds for Improved Approximation Algorithms[1]

Faisal N. Abu-Khzam
Lebanese American University, Beirut, Lebanon
faisal.abukhzam@lau.edu.lb

Cristina Bazgan
PSL, University of Paris-Dauphine, LAMSADE UMR 7243, France;
and Institut Universitaire de France
bazgan@lamsade.dauphine.fr

Morgan Chopin
Institut für Optimierung und Operations Research, Universität Ulm, Germany
morgan.chopin@uni-ulm.de

Henning Fernau
Fachbereich 4, Abteilung Informatikwissenschaften, Universität Trier,
Germany
fernau@uni-trier.de

Abstract

Kernelization algorithms in the context of Parameterized Complexity are often based on a combination of reduction rules and combinatorial insights. We will expose in this paper a similar strategy for obtaining polynomial-time approximation algorithms. Our method features the use of approximation-preserving reductions, akin to the notion of parameterized reductions. We exemplify this method to obtain the currently best approximation algorithms for Harmless Set, Differential and Multiple Nonblocker, all of them can be considered in the context of securing networks or information propagation.

Keywords: Reduction rules, maximization problems, polynomial-time approximation, domination problems

1 Introduction

It is well-known that most interesting combinatorial problems are hard from a computational point of view. More technically speaking, they mostly turn out to be NP-hard. As many of these combinatorial problems have some importance for practical applications, several techniques have been developed to deal with them. From a more mathematical angle, the two most interesting and widespread approaches are (polynomial-time) approximation and fixed-parameter algorithms. Both areas have developed their own set of tools over the years. For instance, methods related to Linear Programming are prominent in the area of Approximation Algorithms [3]. Conversely, data reduction rules are the

[1]This work is supported by the bilateral research cooperation CEDRE between France and Lebanon (grant number 30885TM). An extended abstract will be presented at ISAAC 2014.

method of choice to obtain kernelization results, which is central to Parameterized Algorithms [22]. Another essential ingredient to kernelization algorithms is a collection of combinatorial insights to the specific problem, often (already) supplied by mathematicians working in Combinatorics. It is quite natural to try to employ certain tools from one area to the other one. For example, the title of the paper [32] nicely indicates the intended use of Linear Programming to obtain FPT algorithms. In this paper, we take the opposite approach and show how to use data reduction rules and (constructive) combinatorial insights to obtain approximation algorithms, in particular for maximization problems. Notice that data reduction rules are often used in heuristic approaches, well-established in practical implementations. So, our approach also brings the often more theoretical findings closer to practice.

For the purpose of illustrating our method, we will mainly deal with maximization problems that are obtained from domination-type graph problems. We first describe these problems, using standard graph-theoretic terminology.

Let $G = (V, E)$ be an undirected graph and $D \subseteq V$.

1. D is called a *dominating set* if, for all $x \in V \setminus D$, there is a $y \in D \cap N(x)$. $V \setminus D$ is known as an *enclaveless set* [34] or as a *nonblocker set* [21].

2. D is called a *total dominating set* if, for all $x \in V$, there is a $y \in D \cap N(x)$. $V \setminus D$ has been introduced as a *harmless set* or *robust set* (with unaminity thresholds) in [6].

3. If D can be partitioned as $D = D_1 \cup D_2$ such that, for all $x \in V \setminus D$, there is a $y \in D_2 \cap N(x)$, then (D_2, D_1) defines a *Roman domination function* $f_{D_1,D_2} : V \to \{0, 1, 2\}$ such that $f_{D_1,D_2}(V) = 2|D_2| + |D_1|$. According to [10], $D_0 := V \setminus (D_1 \cup D_2)$ is also known as the *differential (set)* of a graph (as introduced in [30]) if $f_{D_1,D_2}(V)$ is smallest possible.

4. If for all $x \in V \setminus D$, there are k elements in $D \cap N(x)$, then D is a *k-dominating set*, see [17, 20, 25]. We will call $V \setminus D$ a *k-nonblocker set*.

The maximization problems derived from these four definitions are: Nonblocker, Harmless Set, Differential, and k-Nonblocker. Actually, Nonblocker has been looked into by the approximation algorithm community quite a lot in recent years [1, 18, 33], where it is known as the Maximum Star Forest problem. Although these problems are all better known from the minimization perspective, there is a good reason to study them in this complementary way: All of these minimization problems do not possess constant-factor approximations under reasonable complexity assumptions (the reduction shown in [19] for (Total) Dominating Set starts from Set Cover), while the complementary problems can be treated in this favorable way. For Roman Domination, observe that the reduction shown in [24] works from Set Cover, so that again (basically) the same lower bounds follow. This move is related to *differential approximation* [2]. Notice that this comes along with similar properties from the perspective of Parameterized Complexity: While natural parameterizations of the minimizations lead to W[2]-hard problems [22, 24], the

natural parameterizations of the maximization counterparts are fixed-parameter tractable. However, as this is more customary as a combinatorial entity, let us refer (as usual) by $\gamma(G)$ to the size of the smallest dominating set of G, by $\gamma_t(G)$ to the size of the smallest total dominating set, by $\gamma_R(G)$ to the Roman domination number of G, i.e., the smallest value of a Roman domination function of G, and by $\gamma_k(G)$ to the size of the smallest k-dominating set of G.

Some graph-theoretic notations Let $G = (V, E)$ be a simple undirected graph. We denote by $N(x)$ the set of neighbors of vertex x; the cardinality of $N(x)$ is the *degree* of x. A vertex of degree zero is known as an *isolated vertex*, and a vertex of degree one as a *leaf*. The number of vertices of a graph is called its *order*. Given $U \subseteq V$, $G[U]$ denotes the subgraph induced by U. A repetition-free sequence $x_1, \ldots, x_k$ of vertices is a *path* in G (of *length* $k-1$) if $x_i x_{i+1} \in E$ for $i = 1, \ldots, k-1$. A *chain* is an induced path whose interior vertices are of degree two in G. The *diameter* of G is the greatest length of a shortest path in G.

Main Results. We introduce a notion of approximation-preserving reductions analogous to parameter-preserving reductions known in Parameterized Complexity in order to obtain new approximation algorithms. We introduce a general methodology to obtain constant-factor approximations for various problems. For instance, along with an algorithmic version of the upper bound obtained in [29] on the size of a total dominating set, we present a factor-two approximation algorithm for Harmless Set, beating the previously known factor of three [6]. Moreover, we are deriving a factor-$\frac{11}{3}$ approximation algorithm for Differential, which was set up as an open problem in [9], where this approximability question could be only settled for bounded-degree graphs; our approach also improves on the factor-4 approximation exhibited in [8]. However, as in [9] APX-completeness was shown for the degree-bounded case, nothing better than constant-factor approximations can be expected for general graphs. Finally, we present constant-factor approximation algorithms for k-Nonblocker.

Organization of the paper Section 2 explains the use of reduction rules within maximization problems. It also exhibits the general method. Sections 3 and 4 show how to employ our general method to one specific problem in a non-trivial way. Sections 5 and 6 show that the same method can be also applied to other problems. We conclude with discussing further research directions.

2 Approximation preserving reductions for maximization problems

Specializing standard terminology from [3], we can express the following. A maximization problem $\mathcal{P}$ can be specified by a triple $(I_{\mathcal{P}}, \mathrm{SOL}_{\mathcal{P}}, m_{\mathcal{P}})$, where

1. $I_{\mathcal{P}}$ is the set of input instances of $\mathcal{P}$;

2. $\mathrm{SOL}_{\mathcal{P}}$ is a function that associates to $x \in I_{\mathcal{P}}$ the set $\mathrm{SOL}_{\mathcal{P}}(x)$ of feasible solutions of x;

3. $m_{\mathcal{P}}$ provides on (x, y), where $x \in I_{\mathcal{P}}$ and $y \in \mathrm{SOL}_{\mathcal{P}}(x)$, a positive integer which is the value of the solution y.

An optimum solution y^* to x satisfies: (i) $y^* \in \mathrm{SOL}_{\mathcal{P}}(x)$, and (ii) $m_{\mathcal{P}}(y^*) = \max\{m_{\mathcal{P}}(y) \mid y \in \mathrm{SOL}_{\mathcal{P}}(x)\}$. The value $m_{\mathcal{P}}(y^*)$ is also referred to as $m^*_{\mathcal{P}}(x)$ for brevity.

Given a maximization problem $\mathcal{P}$, a *factor-α approximation*, $\alpha \geq 1$, associates to each $x \in I_{\mathcal{P}}$ some $y \in \mathrm{SOL}_{\mathcal{P}}(x)$ such that $\alpha \cdot m_{\mathcal{P}}(x, y) \geq m^*_{\mathcal{P}}(x)$. A solution $y \in \mathrm{SOL}_{\mathcal{P}}(x)$ satisfying $\alpha \cdot m_{\mathcal{P}}(x, y) \geq m^*_{\mathcal{P}}(x)$ is also called an *α-approximate solution* for x.

We are now going to present a first key notion for this paper.

Definition 1 *An* α-preserving reduction*, with $\alpha \geq 1$, is a pair of mappings $\mathtt{inst}_{\mathcal{P}} : I_{\mathcal{P}} \to I_{\mathcal{P}}$ and $\mathtt{sol}_{\mathcal{P}}$ which, given $y' \in \mathrm{SOL}_{\mathcal{P}}(\mathtt{inst}_{\mathcal{P}}(x))$, produces some $y \in \mathrm{SOL}_{\mathcal{P}}(x)$ such that there are constants $a, b \geq 0$ satistying $a \leq \alpha \cdot b$ and the following inequalities:*

1. $m^*_{\mathcal{P}}(\mathtt{inst}_{\mathcal{P}}(x)) + a \geq m^*_{\mathcal{P}}(x)$,

2. *for each $y' \in \mathrm{SOL}_{\mathcal{P}}(\mathtt{inst}_{\mathcal{P}}(x))$, the corresponding solution $y = \mathtt{sol}_{\mathcal{P}}(y')$ satisfies: $m_{\mathcal{P}}(\mathtt{inst}_{\mathcal{P}}(x), y') + b \leq m_{\mathcal{P}}(x, y)$.*

When referring to this definition, we mostly explicitly specify the constants a and b for ease of verification. An important trivial example is given by a pair of identity mappings that are α-preserving for any $\alpha \geq 1$. Notice that a similar notion has been introduced, or implicitly used, in the context of minimization problems in [15, 16, 26].

Theorem 1 *Let $\mathcal{P} = (I_{\mathcal{P}}, \mathrm{SOL}_{\mathcal{P}}, m_{\mathcal{P}})$ be some maximization problem. If the pair $(\mathtt{inst}_{\mathcal{P}}, \mathtt{sol}_{\mathcal{P}})$ describes an α-preserving reduction and if, given some instance x, $y' \in \mathrm{SOL}_{\mathcal{P}}(\mathtt{inst}_{\mathcal{P}}(x))$ is an α-approximate solution for $\mathtt{inst}_{\mathcal{P}}(x)$, then $y = \mathtt{sol}_{\mathcal{P}}(y')$ is an α-approximate solution for x.*

Proof. We have to prove that $\alpha \cdot m_{\mathcal{P}}(x, y) \geq m^*_{\mathcal{P}}(x)$. Now,

$$\frac{m^*_{\mathcal{P}}(x)}{m_{\mathcal{P}}(x, y)} \leq \frac{m^*_{\mathcal{P}}(\mathtt{inst}_{\mathcal{P}}(x)) + a}{m_{\mathcal{P}}(\mathtt{inst}_{\mathcal{P}}(x), y') + b} \leq \frac{\alpha m_{\mathcal{P}}(\mathtt{inst}_{\mathcal{P}}(x), y') + \alpha b}{m_{\mathcal{P}}(\mathtt{inst}_{\mathcal{P}}(x), y') + b} = \alpha$$

as required. QED.

This shows that an α-preserving reduction leads to a special AP-reduction as defined in [3]. But there, these reductions were mainly used to prove hardness results, as it is also the case of [26] that we already mentioned. However, we use this notion to obtain approximation algorithms.

The notion of an α-preserving reduction was coined following the successful example of kernelization reductions known from Parameterized Complexity [22].

One of the nice features of those is that they are usually compiled from simpler rules that are often based on some applicability conditions. In the following, we describe that this also works out for approximation. We need two further notions to make this precise.

We call an α-preserving reduction $(\mathtt{inst}_{\mathcal{P}}, \mathtt{sol}_{\mathcal{P}})$ *strict* if $|\,\mathtt{inst}_{\mathcal{P}}(x)| < |x|$ for all $x \in I_{\mathcal{P}}$, and it is called *polynomial-time computable* if the two mappings comprising the reduction can be computed in polynomial time.

The following lemma is relatively straightforward to prove. Yet, it contains an important message: reduction rules can be composed so that the composition

Lemma 2 *If* $(\mathtt{inst}_{\mathcal{P}}, \mathtt{sol}_{\mathcal{P}})$ *and* $(\mathtt{inst}'_{\mathcal{P}}, \mathtt{sol}'_{\mathcal{P}})$ *are two* α*-preserving reductions, then the composition* $(i, s) := (\mathtt{inst}_{\mathcal{P}} \circ \mathtt{inst}'_{\mathcal{P}}, \mathtt{sol}'_{\mathcal{P}} \circ \mathtt{sol}_{\mathcal{P}})$ *is also an* α*-preserving reduction. If both* $(\mathtt{inst}_{\mathcal{P}}, \mathtt{sol}_{\mathcal{P}})$ *and* $(\mathtt{inst}'_{\mathcal{P}}, \mathtt{sol}'_{\mathcal{P}})$ *are strict (polynomial-time computable, resp.), then the composition* (i, s) *is strict (polynomial-time computable, resp.).*

Proof. Consider a situation described as in the lemma, where the pair of numbers (a, b) shows that $(\mathtt{inst}_{\mathcal{P}}, \mathtt{sol}_{\mathcal{P}})$ is α-preserving and the pair of numbers (a', b') shows that $(\mathtt{inst}'_{\mathcal{P}}, \mathtt{sol}'_{\mathcal{P}})$ is α-preserving. Clearly, if $x \in I_{\mathcal{P}}$, then $\mathtt{inst}_{\mathcal{P}}(x) \in I_{\mathcal{P}}$ and hence $(\mathtt{inst}_{\mathcal{P}} \circ \mathtt{inst}'_{\mathcal{P}})(x) = \mathtt{inst}'_{\mathcal{P}}(\mathtt{inst}_{\mathcal{P}}(x)) \in I_{\mathcal{P}}$, as well. A similar observation applies to the solutions (in the reversed order). We now prove that the composition $\mathtt{inst}_{\mathcal{P}} \circ \mathtt{inst}'_{\mathcal{P}}$ is α-preserving, testified by the pair of numbers $(a + a', b + b')$.

$$\begin{aligned} m^*_{\mathcal{P}}((\underset{\mathcal{P}}{\mathtt{inst}} \circ \underset{\mathcal{P}}{\mathtt{inst}}')(x)) + (a + a') &= (m^*_{\mathcal{P}}((\underset{\mathcal{P}}{\mathtt{inst}}'(\underset{\mathcal{P}}{\mathtt{inst}}(x))) + a') + a \\ &\geq m^*_{\mathcal{P}}(\underset{\mathcal{P}}{\mathtt{inst}}(x)) + a \\ &\geq m^*_{\mathcal{P}}(x) \end{aligned}$$

The computation for the bounds on the solution is similar and hence omitted. The claim on composability of the strictness and the polynomial-time computability are easy to see. QED.

By a trivial induction argument, the previous lemma generalizes to any finite number of reductions that we like to compose.

Conditional reductions In the realm of Kernelization, reductions are often described in some conditional form:

if condition **then do** action

Our previous considerations apply also for this type of conditioned reductions, apart from the fact that an instance may not change, assuming that the reduction was not applicable, which means that the condition was not true for that instance.

First, we have to make clear what the notions of "strictness" and "polynomial time computations" refer to in the context of reduction rules with conditions.

"Strictness" now means that the input will be shortened if the condition is met, and "polynomial time" means two things: a) the condition can be checked in polynomial time and b) the possibly triggered action can be performed in polynomial time. Moreover, often there is a finite collection of conditioned reductions. These can be combined in quite a natural way into a single conditioned reduction. This is formally described in the following lemma.

Lemma 3 *Assume that, for each* $1 \leq i \leq n$,

$$\textbf{\textit{if}}\ \mathit{condition}_i\ \textbf{\textit{then do}}\ \mathit{action}_i$$

is a conditioned α*-preserving reduction. Then, these can be combined into a single conditioned* α*-preserving reduction* $\langle \mathit{combi\text{-}condition}, \mathit{combi\text{-}action} \rangle$ *as follows:*

$$\textbf{\textit{if}}\ \exists i(\mathit{condition}_i)\ \textbf{\textit{then do}}\ \mathit{perform\ some\ applicable\ action}_i$$

If all original conditioned reductions are strict (polynomial-time computable, resp.), then the combined reduction is strict (polynomial-time computable, resp.).

Now, we can present a general recipe how to obtain a polynomial-time factor-α approximation based on α-preserving reductions. The previous lemma shows that the use of a single reduction in the formulation of the next theorem does not lose any generality.

Theorem 4 *Assume that* $\mathcal{P}$ *is some maximation problem. Suppose that*

$$\textbf{\textit{if}}\ \mathit{condition}(x)\ \textbf{\textit{then do}}\ \mathit{action}(x)$$

is some conditioned α*-preserving, strict, polynomial-time computable reduction. Further assume that there is some polynomial-time computable factor-*α *approximation algorithm* A *for* $\mathcal{P}$*, restricted to instances from* $\{x \in I_{\mathcal{P}} \mid \neg \mathit{condition}(x)\}$*. Then, there is a polynomial-time computable factor-*α *approximation algorithm for all instances.*

Proof. The desired algorithm should work as follows. Given an instance x:

1. As long as possible, some α-preserving reductions are performed. This yields the sequence of instances $x = x_0, x_1, \ldots, x_n$.
2. Then, A is applied to the reduced instance $x' := x_n$.
3. As $y_n := y' := A(x')$ is an α-approximate solution for $x' = x_n$, we can successively construct α-approximate solutions y_{n-1} for $x_{n-1}, \ldots, y_1$ for x_1 and finally $y := y_0$ for $x = x_0$.
4. Return y as approximate solution for x.

As compositions of α-preserving reductions yield α-preserving reductions (maintaining some desirable properties), as shown in Lemma 2, any α-approximate solution to x_n can be turned into an α-approximate solution for x. Hence, all claimed properties directly follow by our previous considerations, apart from the polynomial-time claim. Here, observe as the reductions are strict, $n \leq |x|$, so that the **while**-loop terminates after a polynomial number of steps. QED.

The general strategy that we follow can be sketched as follows:

1. Apply (strict, poly-time computable) α-preserving reduction rules as long as possible.
2. Possibly modify the resulting graph so that it meets some requirements from known combinatorial results on the graph parameter of interest.
3. Compute some solution for the modified graph that satisfies the mentioned combinatorial bounds.
4. Construct from this solution a good approximate solution for the original instance.

In order to illustrate the use of this strategy, let us elaborate on Nonblocker, matching a result from [33]. Actually, conceptually this algorithm is even simpler than the one we present for Harmless Set in particular in the following sections. This goes along the lines of the kernelization result by Dehne *et al.* [21], but kernelization needs no constructive proof of the combinatorial backbone result; the non-constructive proof of [12, 31] is hence sufficient.

1. Delete all isolates. (If the resulting graph is of minimum degree at least two, we are ready to directly apply the algorithm of Nguyen *et al.* [33].) This rule is α-preserving for any $\alpha \geq 1$ (with $a = b = 1$).
2. Merge all leaf neighbors into a single vertex. Again, this rule is α-preserving for any $\alpha \geq 1$ (with $a = b = 0$ for a single merge and hence also for a finite sequence of merges).
3. Delete all leaves but one, which is x. This yields the graph G of order n_G.
4. Create a copy G' of the graph G; call the vertices in the new graph by priming the names of vertices of G. Let H be the graph union of G and G' plus the edge xx'. H is of minimum degree at least two by construction.
5. Take the algorithm of Nguyen *et al.* [33] to obtain a dominating set D_H of H satisfying $|D_H| \leq \frac{2}{5}n_H$. Should the solution D_H contain x or x', it is not hard to modify it to contain the leaf neighbors y or y', instead.
6. Hence, $D_G = V_G \cap D_H$ is a dominating set for G with $|D_G| \leq \frac{2}{5}n_G$. Trivially, $N_G = V_G \setminus D_G$ is a nonblocker solution for G that is $\frac{5}{3}$-approximate.

7. As the merging and deletion reductions are α-preserving for each $\alpha \geq 1$, we can safely undo them and hence obtain a $\frac{5}{3}$-approximate solution for the original graph instance.

Better approximation algorithms for NONBLOCKER have been obtained by Chen *et al.* [18] (with a factor of 1.41) and by Athanassopoulos *et al.* [1] (with a factor of 1.244).

Nguyen *et al.* used a slightly different way to obtain their approximation algorithm. Let us reformulate and sketch the result from [33] within our framework. As a reduction rule, they only remove isolates; these would be put into the nonblocker set anyways. The combinatorial result aimed at is the one exhibited by Blank [12] and (independently) by McCuaig and Shepherd [31] that shows that any graph (with seven exceptional graphs) of order n with minimum degree of at least two has a dominating set with at most $\frac{2}{5}n$ vertices. This result is used by first modifying the graph by deleting all leaves and then interconnecting the leaf neighbors so that the minimum degree two requirement is met. It is then shown that it is possible to construct a nonblocker set for G, given a dominating set D_H satisfying $|D_H| \leq \frac{2}{5}n_H$ for the modified graph H. An essential ingredient is a new proof of the mentioned result from [31] that is in fact a polynomial-time algorithm to compute D_H within H. This result was also used by our version of this algorithm given above.

3 Harmless Set

We are now turning towards HARMLESS SET as the most elaborate example of our methodology. First, we are going to present the combinatorial backbone of our result. Let $S_2(G)$ be the set all vertices of degree two within G.

Theorem 5 *(Lam and Wei [29]) Let G be a graph of order n_G and of minimum degree at least two such that $G[S_2(G)]$ decomposes into K_1- and K_2-components. Then, $\gamma_t(G) \leq n_G/2$.*

The proof of this theorem is non-constructive, as it uses tools from extremal combinatorics. We show now how to obtain a polynomial-time algorithm that actually computes a total dominating set (TDS) D with $|D| \leq n_G/2$ under the assumptions of Theorem 5. Let $G = (V, E)$ be a graph with (*) minimum degree at least two and no three consecutive vertices of degree two. This last condition is obviously equivalent to requiring that $G[S_2(G)]$ decomposes into connected components of the form K_1 or K_2. As connected components can be computed consecutively, we can assume that G is connected.

First, we greedily remove edges, as long as the graph still satisfies (*). A TDS computed for the resulting graph is also a TDS for the original graph. For simplicity, we can hence further assume that no edges from G can be removed without violating (*). This is a technical condition needed for applying some of the Lemmas from [29].

We now differentiate two main cases:

- If $S_2(G)$ is an independent set in G, i.e., $G[S_2(G)]$ has no edges, then we have to differentiate further cases when the shortest cycle in G is of length 3, 4, 5, 6, or larger. In each of the cases, Lam and Wei show how to construct a graph G' smaller than G that also satisfies (*).
- Otherwise, $G[S_2(G)]$ contains a K_2-component. Starting out from such a path of length one, the proof of [29, Lemma 6] shows how to construct a set Q of vertices such that the graph $G' = G[V \setminus Q]$ also satisfies (*) and, moreover, $\gamma_t(G) \leq \gamma_t(G') + \frac{|Q|}{2}$ is satisfied.

As some optimum TDS can be surely easily computed for small graphs, the sketched procedure allows to recursively compute some TDS solution for G. Notice in particular that the proofs of Lam and Wei show how to construct a solution for the calling instance from the one obtained for the called instance (in the recursion). Also, it is shown (as explicitly indicated in the second case above) that the claimed bound on the solution size easily follows by induction.

Hence, we can state the following constructive version of the combinatorial result of Lam and Wei:

Theorem 6 *For a given graph $G = (V, E)$ of order n_G that satisfies (*), one can compute a TDS $D \subseteq V$ with $|D| \leq \frac{n_G}{2}$ in polynomial time.*

Observe that a quick analysis of the sketched algorithm indicates a bound of $O(n_G^8)$ for the running time, as one has to actually verify that there are no short cycles in G to match the case analysis. Supposedly, a complete re-analysis of the combinatorial argument could reveal better algorithms, but for the proof of concept of our methodology, this analysis is sufficient here.

Our approximation algorithm for Harmless Set is based on obtaining a (small enough) TDS in a graph H obtained from the input G after a number of modifications (mainly vertex deletions). In the reduction from G to H, we distinguish between the number of deleted vertices d (to get from G to H) and the number of vertices a added to convert the TDS D_H to D_G.

Theorem 7 *Let G be a graph of order n_G and let H be a graph of order n_H obtained from G by deleting d vertices and possibly adding some edges. Let D_G and D_H be TDS solutions of G and H, respectively, such that $a = |D_G| - |D_H| \leq d$. If $|D_H| \leq c \cdot n_H$ and $d \leq \gamma_t(G)$, then $V(G) \setminus D_G$ is a harmless set of G whose size $n_G - |D_G|$ is within a factor of $(1-c)^{-1}$ from optimum.*

Proof. As $n_H = n_G - d$, $|D_G| = |D_H| + a \leq c(n_G - d) + a = cn_G + (a - cd) \leq cn_G + d - cd = cn_G + (1-c)d \leq cn_G + (1-c)\gamma_t(G)$. Hence, $n_G - |D_G| \geq n_G - cn_G - (1-c)\gamma_t(G) = (1-c)(n_G - \gamma_t(G))$. This immediately yields an approximation factor of $(1-c)^{-1}$. QED.

In the following section, we will present reduction rules that produce a graph G with the property (**) that each vertex of degree bigger than one has at most one leaf neighbor. The surgery that produces a graph H from G as indicated in Theorem 7 includes removing all d leaves and adding edges to ensure that

H has minimum degree of two and satisfies that each component of $H[S_2(H)]$ has diameter at most one. Notice that all leaf neighbors in G belong to some optimum TDS of G without loss of generality. Due to (**), $\gamma_t(G) \geq d$ as required. Moreover, given some TDS solution D_H for H, we can produce a valid TDS solution D_G for G by adding all d leaf neighbors to D_H. Notice that Theorem 7 leads to a factor-2 approximation algorithm for Harmless Set based on Theorem 6.

In the following section, we are going to describe the reduction rules necessary to produce a graph to which we could apply the mentioned combinatorial results.

4 Reduction Rules for Harmless Set

Now, we list α-preserving reductions for Harmless Set. We start with two very simple rules.

Isolate Reduction If there is some isolated vertex, produce the instance $(\{x\}, \emptyset)$ that has trivially no solution. If there is some isolated edge xy, produce that instance $G[V \setminus \{x, y\}]$ from $G = (V, E)$.

For the correctness of this rule, observe that a graph with isolated vertices has no total dominating set at all.

Leaf Reduction If there are two leaf vertices u, v with common neighbor w, then delete u. (It would go into the harmless set.)

Observation 8 *The Isolate Reduction (for edges) and the Leaf Reduction are α-preserving for any $\alpha \geq 1$.*

Proof. The Isolate Reduction is α-preserving by setting $a = b = 0$ in the definition. In other words, endpoints of isolated edges must belong to any TDS solution. The Leaf Reduction is α-preserving by setting $a = b = 1$ in the definition. In other words, w.l.o.g., leaves do not belong to some TDS solution, except when there is a K_2-component in the graph. QED.

Hence from now on, no vertex can have two leaf neighbors.

Actually, we could generalize the Leaf Reduction towards the following rule:

Twin Reduction Recall that vertices u and v are said to be *true twins* if $N[u] = N[v]$ and *false twin* if $N(u) = N(v)$.

- If there are two vertices u and v such that $N[u] = N[v]$, i.e., they form true twins, then delete v (it would go into the harmless set).
- If there are two vertices u and v such that $N(u) = N(v)$, i.e., they form false twins, then delete v (it would go into the harmless set).

As we are not using this rule in some crucial manner in what follows, we present the following result without proof.

Theorem 9 *The Twin Reduction is α-preserving for any $\alpha \geq 1$.*

A chain with one leaf endpoint is a *pendant chain*. A *floating chain* is a chain with two leaves. A *support vertex* is a non-pendant endpoint of a pendant chain. Support vertices may have more than one pendant chain. We shall reduce the length of pendant chains to at most two, based on the following reduction rules. The first one actually generalizes the Isolate Reduction.

Floating Chain Reduction Delete all floating chains.

Observation 10 *The Floating Chain Reduction rule is α-preserving for any $\alpha \geq 1$.*

Proof. G' is obtained from G by deleting a floating chain. For the chain, the numbers $a = b$ can be computed (optimally) in polynomial time; they correspond to the size of optimum solutions for the floating chain component. QED.

Long Chain Reduction Assume that G is a graph that contains a path $x - u - v - w - y$, where u, v, w are three consecutive vertices of degree two, where $|N(y)| \geq 2$. Then, construct the graph G' by

- deleting x, u, v, w and
- connecting y to all vertices in $N(x) \setminus \{u\}$ (without creating double edges).

This corresponds to merging x and y and deleting u, v, w. This Long Chain Reduction resembles the folding rule known for Vertex Cover (in Parameterized Complexity, see [22]).

Theorem 11 *The Long Chain Reduction is α-preserving for any $\alpha \geq 1$.*

Proof. Let G be the original graph and G' the graph obtained from G by deleting the path u, v, w and merging x and y as described by the rule. We show that $a = b = 2$ works out in our case by considering several cases.

(a) Let C be a maximum harmless set (HS) for G. Let us first briefly discuss what happens if $N(x) = \{u\}$. Then, it is not hard to see that an optimum solution C would contain x and w, but not u and v. Merging x and y and deleting u, v, w is now equivalent to deleting the whole pending path $x-u-v-w$. As $w \in C$, it does not dominate y, so that $C' = C \setminus \{x, w\}$ is a valid harmless set for G'.

In the following discussion, we can hence assume that x has at least two neighbors. We now consider cases whether or not $x \in C$ or $y \in C$.

- Assume that $x \in C$ and $y \in C$. Hence, u, w are not dominated neither by x nor by y. As C is maximum, we can assume $|C \cap \{u, v, w\}| = 1$, as $\min\{|N(x)|, |N(y)|\} \geq 2$; hence, if all of u, v and w are in $V \setminus C$, then we can replace w by another neighbor of y and obtain another optimum solution. Then, $C' = C \setminus \{x, u, v, w\}$ is a HS of G', with $|C'| = |C| - 2$.

- Assume that $x \notin C$ and $y \notin C$. First, let us discuss the possibility that $u \notin C$ and $w \notin C$. As C is maximum, the purpose of this is to dominate (i) v and (ii) x and y. To accomplish (i), either $u \notin C$ or $w \notin C$ would suffice. However, as C is maximum, condition (ii) means that $N(x) \setminus C = \{u\}$ and that $N(y) \setminus C = \{w\}$. By our assumptions, $\min\{|N(x)|, |N(y)|\} \geq 2$. Hence, there is a vertex $z \in N(y)$, $z \neq w$. Now, $\tilde{C} = (C \setminus \{z\}) \cup \{w\}$ is also a maximum HS satisfying $\{v, w\} \subseteq \tilde{C}$. From now on, we assume that $|C \cap \{u, v, w\}| = 2$ and that $|((N(x) \cup N(y)) \setminus (\{u, w\} \cup C)| \geq 1$ (in other words, at least one of x and y has a neighbor in $V \setminus C$ other than u and w, respectively). Hence, $C' = C \setminus \{u, v, w\}$ is a HS of G' with $|C'| = |C| - 2$.

- Assume now that $x \in C$ and $y \notin C$. (Clearly, the case that $x \notin C$ and $y \in C$ is symmetric.) As u is not dominated by x, either (i) $\{u, v\} \subseteq V \setminus C$ or (ii) $\{v, w\} \subseteq V \setminus C$. In case (i), x is dominated by u, but y must (still) be dominated by some vertex from $N(y) \setminus \{w\}$. In case (ii), symmmetrically y is dominated by w, but x must be dominated by some vertex from $N(x) \setminus \{u\}$. In both cases, $\tilde{C} = (C \setminus \{x\}) \cup \{v\}$ is another maximum harmless set of G. This leads us back to the previous item (i.e., $|C'| = |C| - 2$.)

Summarizing, we have shown that from C we can construct a harmless set C' for G' with $|C'| = |C| - 2$.

(b) Conversely, assume C' is some harmless set for G'. We distinguish two cases:

- Assume that $y \in C'$. Then, y is dominated by some z in its neighborhood (in G'). We consider two cases according to the situation in G. (i) If $z \in N(x)$, then $C = C' \cup \{x, u\}$ is a HS in G. (ii) If $z \in N(y)$, then $C = C' \cup \{x, w\}$ is a HS in G. In both cases, $|C| = |C'| + 2$.

- If $y \notin C$, then again y is dominated by some z in its neighborhood (in G'). We perform the same case distinction as in the previous case: (i) If $z \in N(x)$, then $C = C' \cup \{u, v\}$ is a HS in G. (ii) If $z \in N(y)$, then $C = C' \cup \{v, w\}$ is a HS in G. In both cases, $|C| = |C'| + 2$.

(c) The reasoning from (b) shows that, if C is an optimum solution for G, then C' as obtained in part (a) of this proof is an optimum solution for G'. Namely, assume that there would be a harmless set C^* for G' with $|C^*| > |C'|$. Then, according to (b), we can construct a harmless set of G with $|C^*| + 2 > |C|$ many vertices, contradicting the maximality of C. QED.

Similarly, one sees the correctness of the following rule.

Cycle Chain Reduction If G is a graph that contains a cycle $x-u-v-w-x$, where u, v, w are three consecutive vertices of degree two, then construct the graph G' by deleting u.

Observation 12 *The cycle chain reduction is α-preserving for any $\alpha \geq 1$.*

Proof. An optimum harmless set for G will put exactly two out of the three vertices u, v, w into the harmless set. W.l.o.g., let these be u and v. Conversely, w and x would go into the total dominating set. Also, in the reduced graph, v will be in the harmless set, while x and w will be in the total dominating set. This shows the claim with constants $a = b = 1$. QED.

Finally, we deal with support vertices with multiple pendant chains. Assuming the Long Chain Reduction has been applied, any pendant chain is of length two or less. Accordingly, a support vertex where two of more pendant chains meet does belong to some optimum solution. The following rule makes this idea more precise.

Pendant Chain Reduction Assume that $G = (V, E)$ is a graph that contains two pendant chains with common endpoint v of which at least one path is of length two. Then, construct the graph $G' = (V', E')$ by deleting one of the two pendant chains, keeping one which is of length two.

Theorem 13 *The Pendant Chain Reduction is α-preserving for any $\alpha \geq 1$.*

Proof. Let $v-x-y$ and (a) $v-z-t$ (or (b) just $v-t$) be two pendant paths of G. Then y belongs to some maximum harmless set C of G while x belongs to $V(G) \setminus C$. Similarly, z (if existent) belongs to some maximum harmless set C of G while t belongs to $V(G) \setminus C$. It follows that the definition of α-preserving reduction can be applied with $a = b = 1$. Notice that, because we keep $v-x-y$, neither v nor x will belong to any harmless set solution for the reduced graph. Hence, adding t to the harmless set solution of the reduced graph is always possible, resulting in a valid harmless set for the original graph. QED.

We are now in the position to apply Theorem 7.

Observation 14 *Assume the graph $G = (V, E)$ is reduced according to the reduction rules described so far. Hence, G satisfies the following properties:*

- *G contains no chain of three vertices of degree two.*
- *By the Leaf Reduction rule, any vertex has at most one leaf neighbor.*

Let G' be a graph isomorphic to G so that each vertex v of G corresponds to a vertex v' of G', under the assumed isomorphism $f : V(G) \longrightarrow V(G')$. We construct a graph H obtained from the disjoint union of G and G' simply by adding edges between each leaf neighbors vertex v of G with $v' = f(v) \in V(G')$. Then, we remove all leaves.

Due to the application of the Pendant Chain Reduction rule to G (and G'), the addition of edges between corresponding leaf neighbors in G and G' does not introduce induced cycles with more than two consecutive degree-two vertices.

To the resulting graph H, apply Long Chain Reduction as long as possible. Notice that an application of this rule does never decrease degrees, adds two vertices to the solution and removes four vertices of the graph.

This results in a graph H' of order $n_{H'}$ with minimum degree at least two containing no chain of three vertices of degree two. Hence, we can apply the (algorithmic) version of Theorem 5 that returns a TDS $D_{H'}$ for H' with $2|D_{H'}| \geq n_{H'}$. Undoing a certain number of Long Chain Reductions, say, c, that we applied, we obtain a TDS D_H for H with $2|D_H| = 2(|D_{H'}|+2c) \geq n_{H'}+4c = n_H$. By symmetry, we can assume that $|D_H \cap V(G)| \leq |D_H \cap V(G')|$. Now, we add all support vertices to $D_H \cap V(G)$ and further vertices to obtain D_G by the following rules:

- If a support vertex already belongs to D_H, then it could have been dominated via the edge that we introduced. As this interconnects to another support vertex, both already belonged to D_H. We arbitrarily select two neighbors (in G) of these support vertices and put them into D_G. Hence, the mentioned support vertices and the attached leaves are totally dominated.

- If a support vertex x did not already belong to D_H, two cases arise: (a) If it was dominated (in H) via an edge already belonging to G, then we do nothing on top of what we said. (b) If the support vertex x was dominated (in H) by an edge xy not belonging to G, then we must add another neighbor z (in G) of x to D_G. However, as (obviously) the vertex y belonged to D_H and was dominated by a neighbor (in G) in D_H, we add (in total) two vertices x, z for the two support vertices x, y. Seen from the other side, this covers the case of a support vertex that already belonged to D_H but was not dominated via the edge that we introduced.

Altogether, we see that we delete all leaves and introduce at most that many vertices into D_G (in comparison to $D_H \cap V(G)$). By Theorem 7 and since all reduction rules take polynomial time, we obtain:

Theorem 15 Harmless Set *is factor-2 polynomial-time approximable. QED.*

5 The differential of a graph

Let us start with an alternative presentation of this notion. Let $G = (V, E)$ be a graph. For $D_0 \subseteq V$, let

$$\partial(D_0) := \left| \left(\bigcup_{x \in D_0} N(x) \right) \setminus D_0 \right| - |D_0|.$$

$\partial(D_0)$ is called the *differential* of the set D_0, and our aim is to find a vertex set that maximizes this quantity. This maximum quantity is known as the differential of G, written $\partial(G)$. The following combinatorial results are known:

Theorem 16 *[7] Let G be a connected graph of order n.*

- *If $n \geq 3$, then $\partial(G) \geq n/5$.*
- *If G has minimum degree at least two, then $\partial(G) \geq \frac{3n}{11}$, apart from five exceptional graphs, none of them having more than seven vertices.*

It is not hard to turn the first combinatorial result into a kernelization result, yielding a kernel bound of $5k$, where k is the natural parameterization of the DIFFERENTIAL. Along the lines of [6], we can obtain a factor-5 approximation by first computing a spanning tree $T = (V, E_T)$ for G and then computing an optimum differential set D_T in T by dynamic programming, and then observing that D_T is a factor-5 approximation for G. In [8], this result was improved to a kernel whose order is bounded by $4k$. Along those lines, we can also get a factor-4 approximation. However, the second item of Theorem 16 suggests a possible improvement to a factor of $\frac{11}{3}$ if we employ our framework. This is what we are going to endeavor in this section.

First, we have to show that the reduction rules presented in [8] as kernelization rules can be also interpreted as α-preserving rules. We use some non-standard terminology for stating the rules. A *hair* is a sequence of two vertices uv, where u is a leaf and v has degree two. Then, u is also called a *hair leaf*.

We use the following simple notation for a hair uv for reasons of clarity: $u - v - \cdots$.

1. **Leaf Reduction.** If there are two leaves connected to the same vertex, then connect these leaves.
2. **Hair Reduction.** If there are two hairs connected to the same vertex, then remove the two hair leaves.
3. **Leaf-Hear Reduction.** If there is a leaf and a hair connected to the same vertex, then remove the hair leaf.
4. **Long Hair Reduction.** If there is a hair $u - v - \cdots$ connected to a vertex w of degree two, then remove u, v, w.
5. **Neighbor Hair Reduction.** If there is a hair $u - v - \cdots$ connected to a vertex w and another hair $u' - v' - \cdots$ connected to a neighbor w' of w, then remove the edge ww'.

In the reasoning given for the rules in [8], only for the Long Hair Reduction, the natural parameter changes (decreases by one). The argument shows (for the other cases) that even if a set of vertices is produced for the reduced graph that is not a valid solution for the original graph, still another solution can be constructed that is not worse (smaller) than the one that was obtained, so that

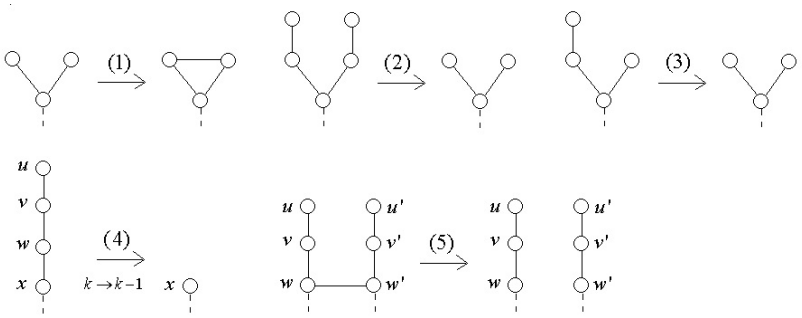

Figure 1: Reduction rules for DIFFERENTIAL.

approximation factors are clearly preserved. The Long Hair Reduction can be seen to be α-preserving when setting $a = b = 1$. Hence, we can summarize:

Observation 17 *The previous five reductions are α-preserving for any $\alpha \geq 1$.*

Lemma 18 *[8] Let $G = (V, E)$ be a graph where none of the previous five reductions applies. Then, G has the following properties:*

(1) *To each vertex, at most one leaf or one hair is attached, but not both together.*

(2) *If we remove all leaves and all hairs from G, then the remaining graph $\bar{G} = (\bar{V}, \bar{E})$, henceforth called* nucleus*, has minimum degree of at least two.*

(3) *If a hair is attached to a vertex u in the nucleus, then no hair is attached to any neighbor of u within the nucleus.*

Notice that the properties listed in Lemma 18 ensure that when obtaining the nucleus H from the reduced graph G by deleting d vertices, $d \leq \gamma_R(G)$ is verified. In order to verify that a sufficiently big solution for the nucleus can be found in polynomial time, observe the proof strategy of [7]: There, the differential of a graph is modeled by so-called big star packings. It is possible to start with a greedily obtained big star packing and then further modify the solution, using the local (and hence easy-to-check) criteria exhibited in various lemmas of that paper, up to the point when no further improvements are possible. The big star packing obtained in this way corresponds to a differential set D with $\partial(D) \geq \frac{3}{11}n$, where n is the order of the graph.

As the proof in [7] uses extremal combinatorial arguments, it is (at least at first glance) non-constructive. Let us give some more details of the algorithm that is hidden within these combinatorial arguments in the following.

The greedy selection of a big star packing. We will (from now on) work on a mixed graph (i.e., a graph that has both directed and undirected edges; directed edges are also called arcs) such that each vertex has at most one outgoing arc, and no vertex with an incoming arc has an outgoing arc. We will call a vertex incident to some directed arc *marked*. As we start with an undirected graph $G = (V, E)$, at the beginning all vertices are unmarked. We proceed as follows:

As long as possible:

- Pick some unmarked vertex x with at least two unmarked neighbors.
- Direct all edges connecting x to any unmarked neighbor towards x.

Now, consider the set D of vertices to which some arcs point to, and let $B(D)$ denote the remaining marked vertices. Clearly, $\left(\bigcup_{x\in D} N(x)\right)\setminus D = B(D)$. Hence, $\partial(D) = |B(D)| - |D|$. Moreover, due to the directions of the edges, we can view each $x \in D$ as the center of a star to which at least two arcs (rays) are pointing. So, we have defined a collection $\mathcal{S}(D)$ of stars that can be viewed as a star packing. As each star has at least two rays, we called them *big stars*. Due to our greedy approach, we hence arrive at $\mathcal{S}(D)$ as being a maximal big star packing. Moreover, $|B(D)|$ is also the number of directed edges (or rays) in total. Let $C(D) := V \setminus (B(D) \cup D)$.

By definition of the partition $(D, B(D), C(D))$ of V we find:

Observation 19 *No edge connects vertices from D with vertices from $C(D)$.*

As we obtain a maximal big star packing, we conclude:

Observation 20 *The induced graph $G[C(D)]$ is undirected and decomposes into K_1- and K_2-components.*

First local improvement. We are now going to improve the solution found so far.

As long as possible:

- Pick some vertex x from $B(D)$ that has two or more neighbors in $C(D)$.
- Let $y \in D$ be such that the edge xy is directed towards y.
- Replace the arc from x to y by an undirected edge again.
- If there is now (only) one arc zy directed to y, remove y from D and render zy an undirected edge again. (This will increase the number of unmarked vertices.)
- Direct all edges that connect x to some umarked vertex towards x and put x into D.

If no further improvements are possible, one might want to ensure that the (new) big star packing $\mathcal{S}(D)$ is still maximal. If not, obvious further improvements are possible. However, after a finite number of steps, this will end.

The new set D (and the related star packing $\mathcal{S}(D)$ that can be read off from the directed edges) satisfies Observation 20 and:

Observation 21 *Every $x \in B(D)$ has at most one neighbor in $C(D)$.*

Second local improvement. By a procedure similar to the previous case, we can create new stars if some $x \in B(D)$ is part of a star with at least three rays and neighbor of some K_2-component in $G[C(D)]$. Leaving out details in this case, we can observe for the (new) differential set D:

Observation 22 *If $x \in B(D)$ is neighbor of some K_2-component in $C(D)$, then it belongs to some star with at least three rays.*

Some simple computations (as undertaken in [7]) show that the set D satisfies our desired bound, i.e., $|D| \geq \frac{3}{11}|V|$, if the packing $\mathcal{S}(D)$ only contains stars with at least four rays. Anyways, it could well be that the (valid) differential set D satisfies the bound and we can stop here.

Further local improvements on smaller stars $K_{1,2}$. If not, then we have to make further local improvements on these smaller stars, considering them in groups. The (relatively messy) details can be found in Lemmas 3.10 through 3.17 in [7], but this should make clear that finally we can obtain a sufficiently big differential in polynomial time.

Having obtained such differential set for the nucleus of a graph, this solution can be easily lifted to a solution of the reduced graph; Theorem 23 allows us to conclude Theorem 24.
We are going to use the idea of computing a sufficiently big solution for the nucleus, based on the following variant of Theorem 7.

Theorem 23 *Let G be a graph of order n_G and let H be a graph of order n_H obtained from G by deleting d vertices. Let $D_G = D_{G,1} \cup D_{G,2}$ and $D_H = D_{H,1} \cup D_{H,2}$ be Roman DS solutions of G and H, respectively, such that $D_{H,2} = D_{G,2}$ and $a = |D_{G,1}| - |D_{H,1}| \leq d$. If $|D_{H,1}| + 2|D_{H,2}| \leq c \cdot n_H$ and $d \leq \gamma_R(G)$, then $\partial(V(G) \setminus D_G) = n_G - 2|D_{G,2}| - |D_{G,1}|$ is within a factor of $(1-c)^{-1}$ from optimum.*

Proof. As $n_H = n_G - d$, $|D_{G,1}| + 2|D_{G,2}| = |D_{H,1}| + a + 2|D_{H,2}| \leq c(n_G - d) + a = cn_G + (a - cd) \leq cn_G + d - cd = cn_G + (1-c)d \leq cn_G + (1-c)\gamma_R(G)$. Hence, $n_G - 2|D_{G,2}| - |D_{G,1}| \geq n_G - cn_G - (1-c)\gamma_t(G) = (1-c)(n_G - \gamma_R(G)) = (1-c)\partial(G)$. This immediately yields an approximation factor of $(1-c)^{-1}$.QED.

We can turn the (non-constructive) combinatorial reasoning of [7] into a polynomial-time algorithm, which allows us to conclude with our framework:

Theorem 24 DIFFERENTIAL *is factor-$\frac{11}{3}$ polynomial-time approximable.*

6 Multiple Nonblocker sets

We are first going to explain why neither some nice approximation algorithm nor some FPT algorithm (with the standard parameterization) yields useful results. We shall assume $k > 1$ in this section.

Theorem 25 *k-Dominating Set, $k > 1$ cannot be better approximated than Dominating Set. Likewise, the (standard) parameterized version is W[2]-hard.*

Proof. Namely, given an instance G of Dominating Set, we introduce (in total) k copies of each vertex, say, $v[1], \ldots, v[k]$ of vertex v, and introduce a $K_{k,k}$ in $\{u[1], \ldots, u[k]\} \cup \{v[1], \ldots, v[k]\}$ whenever there is an edge uv in G. Then, the new graph has a k-dominating set of size kt if and only if the original graph G has a dominating set of size t.

We consider now a combinatorial upper bound on the size of some feasible solution of the minimization problem.

Theorem 26 ([20]) *Let G be a graph of order n_G and a minimum degree at least k. Then $\gamma_k(G) \leq \frac{k}{k+1} n_G$.*

The known non-constructive proof can be turned into a polynomial-time algorithm obtaining the following result.

Theorem 27 *For a given graph G of order n_G and minimum degree at least k, one can compute a k-dominating set D with $|D| \leq \frac{k}{k+1} n_G$ in polynomial time.*

Proof. First, we greedily remove edges between vertices of degree greater than k obtaining a graph G' of minimum degree (exactly) k. Let $S = \{v \in V : d(v) > k\}$. By construction, S is an independent set in G'. We build a maximal independent set T that contains S. Then $V \setminus T$ is a k-dominating set.

If $|V \setminus T| \leq kn_G/(k+1)$, then $D := V \setminus T$ is also a k-dominating set in the supergraph G of G'. Otherwise, while $|V \setminus T| > kn_G/(k+1)$, construct a maximal independent set T' of $G[V \setminus T]$ and set $T = T'$. We show in the following that the algorithm terminates.

Let $r = |T|$. When $|V \setminus T| > kn_G/(k+1)$, we get $n_G = r + |V \setminus T| > r + kn_G/(k+1) = r + k(r + |V \setminus T|)/(k+1)$, thus $|V \setminus T| > kr$. Since T is a maximal independent set (and hence a dominating set) and every element of $V \setminus T$ is of degree k in G', every vertex of $V \setminus T$ has degree at most $k-1$ in $G'[V \setminus T]$. It follows that any maximal independent set of $G'[V \setminus T]$ contains at least $r+1$ vertices (otherwise, $|V \setminus T| \leq r + r(k-1) = rk$). Compute any maximal independent set T' of $G'[V \setminus T]$. Now $|T'| > r = |T|$ and $V \setminus T'$ is a k-dominating set that is smaller than $V \setminus T$; namely, because the minimum degree is at least k, all elements of any independent set are k-dominated by its complement. QED.

Theorem 28 *Let G be a graph of order n_G and let H be a graph of order n_H obtained from G by deleting d vertices and adding $2k$ new vertices, with $d > k$. Let D_G and D_H be k-dominating set solutions of G and H such that $a = |D_G| - |D_H| = d - k$. If $|D_H| \leq c \cdot n_H$ and $d \leq \gamma_k(G)$, then $V(G) \setminus D_G$ is a k-nonblocker of G whose size $n_G - |D_G|$ is within a factor of $(1-c)^{-1}$ from optimum (modulo an additive constant less than k).*

Proof. As $n_H = n_G - d + 2k$, $|D_G| = |D_H| + a \leq c(n_G - d + 2k) + d - k \leq cn_G + (1-c)d + 2ck - k \leq cn_G + (1-c)\gamma_k(G) + k(2c-1)$. Hence, $n_G - |D_G| \geq n_G - cn_G - (1-c)\gamma_k(G) - k(2c-1) = (1-c)(n_G - \gamma_k(G)) - k(2c-1)$. This immediately yields an approximation factor of $(1-c)^{-1}$ (modulo the additive constant $k(2c-1) \leq \frac{k(k-1)}{k+1} < k$). QED.

In the rest of this section, we present reduction rules that produce a graph G with minimum degree at least k. Our reduction rules mainly deal with vertices of degree $k-1$ or less. Each such vertex must be in any k-dominating set. We shall refer to such vertices by *low-degree* vertices in the sequel.

Low-Degree Vertex Deletion Reduction If a low-degree vertex v has only low-degree neighbors, then delete v. If there is a vertex u with $k+1$ low-degree neighbors, then delete one neighbor of u.

Observation 29 *The Low-Degree Vertex Deletion Reduction is α-preserving for any $\alpha \geq 1$.*

Proof. The soundness of Low-Degree Vertex Deletion is rather straightforward. A low-degree vertex that is not a neighbor of a high-degree vertex can be placed (safely) in any k-dominating set. If the number of low-degree neighbors of a vertex u is $t > k$, then we can safely delete $t - k$ such neighbors and place them in the k-dominating set. We keep k neighbors to make sure any subsequent solution places u in the nonblocker set. This reduction is α-preserving with constants $a = b = 0$. QED.

Low-Degree Merging Reduction Let G be an instance of k-NONBLOCKER that has been subject to the Low-Degree Vertex Deletion Reduction rule. Then we add a complete bipartite graph $K_{k,k}$ with new vertices $u_1, \ldots, u_k$, $v_1, \ldots, v_k$. For every high-degree vertex $v \in V$ having q low-degree neighbors $w_1, \ldots, w_q$, with $q \leq k$, delete $w_1, \ldots, w_q$, and connect v to $v_1, \ldots, v_q$.

Observation 30 *The Low-Degree Merging Reduction rule is α-preserving for any $\alpha \geq 1$.*

Proof. We are going to verify the definition of α-preserving reductions; to this end, we show that $a = b = k$ works out in our case. Let $G = (V, E)$ be the original graph and $G' = (V', E')$ the graph obtained from G by deleting $w_1, \ldots, w_q$, and connect v to $v_1, \ldots, v_q$.

(a) Let C be a maximum k-nonblocker for G. C does not contain $w_1, \ldots, w_q$ since $w_1, \ldots, w_q$ are part of any k-dominating set. Consider $C' = C \cup \{u_1, \ldots, u_k\}$. C' is a maximum k-nonblocker set for G' of size $|C'| = |C| + k$.

(b) Consider now the converse. Let C' be some k-nonblocker set for G'. We can suppose that C' contains $u_1, \ldots, u_k$, otherwise we remove $v_1, \ldots, v_k$ and add $u_1, \ldots, u_k$. Thus $C = C' \setminus \{u_1, \ldots, u_k\}$ is a k-nonblocker set for G of size $|C| = |C'| - k$. QED. The reductions above take polynomial time, so that Theorem 28 allows us to conclude:

Theorem 31 *For any instance x of k-*Nonblocker*, one can compute in polynomial time a k-nonblocker set S with $|S| \geq \frac{m^*(x)}{k+1} - k$.*

Combinations with the previous section as indicated in the definitions of [27] should be possible. We leave this for future research, similar to variants like Liar's Domination; see [11] and the literature quoted therein.

7 Conclusions

We presented a framework for obtaining approximation algorithms for maximization problems, inspired by similar reasonings for obtaining kernelization results. We see five major directions from this approach:

- Paraphrasing [23], we might say that not only FPT, but also polynomial-time maximization is *P-time extremal structure*. This should inspire mathematicians working in graph theory (and other areas of combinatorics) to work out useful combinatorial bounds on different graph parameters. We started on domination-type parameters, and this might be a first venue of continuation, for example, along the lines sketched in [13, 14, 28].

- Conversely, approximation algorithms that stay within the combinatorial grounds of their problem tend to reveal (combinatorial) insights into the problem that might get lost when moving for instance into the area of Mathematical Programming.

- The notion of α-preserving reduction is similar to the local ratio techniques [4] that allowed to re-interpret many (e.g., primal-dual) approximation algorithms (for minimization problems) in a purely combinatorial fashion; see [5]. We see some hope for similar developments using α-preserving reduction for maximization problems.

- The fact that α-preserving reductions are inspired by FPT techniques should allow to adapt these notions for obtaining new and faster parameterized approximation algorithms.

- Reductions are often close to practical heuristics and hence allow for fast implementations.

Acknowledgements. We are grateful for discussing this paper at the Bertinoro Workshop on Parameterized Approximation in May 2014.

References

[1] S. Athanassopoulos, I. Caragiannis, C. Kaklamanis, and M. Kyropoulou. An improved approximation bound for spanning star forest and color saving. In R. Královic and D. Niwinski, editors, *Mathematical Foundations of Computer Science 2009, 34th International Symposium, MFCS*, volume 5734 of *LNCS*, pages 90–101. Springer, 2009.

[2] G. Ausiello, C. Bazgan, M. Demange, and V. Th. Paschos. Completeness in differential approximation classes. *International Journal of Foundations of Computer Science*, 16(6):1267–1295, 2005.

[3] G. Ausiello, P. Creczenzi, G. Gambosi, V. Kann, A. Marchetti-Spaccamela, and M. Protasi. *Complexity and Approximation; Combinatorial Optimization Problems and Their Approximability Properties*. Springer, 1999.

[4] R. Bar-Yehuda, K. Bendel, A. Freund, and D. Rawitz. Local ratio: a unified framework for approximation algorithms. *ACM Surveys*, 36(4):422–463, 2004.

[5] R. Bar-Yehuda and D. Rawitz. On the equivalence between the primal-dual schema and the local-ratio technique. *SIAM Journal of Discrete Mathematics*, 19:762–797, 2005.

[6] C. Bazgan and M. Chopin. The robust set problem: Parameterized complexity and approximation. In B. Rovan, V. Sassone, and P. Widmayer, editors, *Mathematical Foundations of Computer Science 2012 – 37th International Symposium, MFCS*, volume 7464 of *LNCS*, pages 136–147. Springer, 2012.

[7] S. Bermudo and H. Fernau. Lower bounds on the differential of a graph. *Discrete Mathematics*, 312:3236–3250, 2012.

[8] S. Bermudo and H. Fernau. Combinatorics for smaller kernels: The differential of a graph. *Submitted*, 2013.

[9] S. Bermudo and H. Fernau. Computing the differential of a graph: hardness, approximability and exact algorithms. *Discrete Applied Mathematics*, 165:69–82, 2014.

[10] S. Bermudo, H. Fernau, and J. M. Sigarreta. The differential and the Roman domination number of a graph. *Applicable Analysis and Discrete Mathematics*, 8:155–171, 2014.

[11] A. Bishnu, A. Ghosh, and S. Paul. Parameterized complexity of k-tuple and liar's domination. *CoRR*, abs/1309.5461, 2013.

[12] M. Blank. An estimate of the external stability number of a graph without suspended vertices (in Russian). *Prikl. Math. i Programmirovanie Vyp.*, 10:3–11, 1973.

[13] M. Borowiecki and D. Michalak. Generalized independence and domination in graphs. *Discrete Mathematics*, 191:51–56, 1998.

[14] A. Bouchou, M. Blidia, and M. Chellali. Relations between the Roman k-domination and Roman domination numbers in graphs. *Discrete Math., Alg. and Appl.*, 6(3), 2014.

[15] L. Brankovic and H. Fernau. Parameterized approximation algorithms for HITTING SET. In R. Solis-Oba and G. Persiano, editors, *Approximation and Online Algorithms — 9th International Workshop, WAOA 2011*, volume 7164 of *LNCS*, pages 63–76. Springer, 2012.

[16] L. Brankovic and H. Fernau. A novel parameterised approximation algorithm for MINIMUM VERTEX COVER. *Theoretical Computer Science*, 511:85–108, 2013.

[17] Y. Caro and Y. Roditty. A note on the k-domination number of a graph. *International Journal of Mathematics and Mathematical Sciences*, 13(1):205–206, 1990.

[18] N. Chen, R. Engelberg, C. T. Nguyen, P. Raghavendra, A. Rudra, and G. Singh. Improved approximation algorithms for the spanning star forest problem. 65(3):498–516, 2013.

[19] M. Chlebík and J. Chlebíková. Approximation hardness of dominating set problems in bounded degree graphs. *Information and Computation*, 206:1264–1275, 2008.

[20] E. J. Cockayne, B. Gamble, and B. Shepherd. An upper bound for the k-domination number of a graph. *Journal of Graph Theory*, 9:533–534, 1985.

[21] F. Dehne, M. Fellows, H. Fernau, E. Prieto, and F. Rosamond. NONBLOCKER: parameterized algorithmics for MINIMUM DOMINATING SET. In J. Štuller, J. Wiedermann, G. Tel, J. Pokorný, and M. Bielikova, editors, *Software Seminar SOFSEM*, volume 3831 of *LNCS*, pages 237–245. Springer, 2006.

[22] R. G. Downey and M. R. Fellows. *Fundamentals of Parameterized Complexity*. Texts in Computer Science. Springer, 2013.

[23] V. Estivill-Castro, M. R. Fellows, M. A. Langston, and F. A. Rosamond. FPT is P-time extremal structure I. In H. Broersma, M. Johnson, and S. Szeider, editors, *Algorithms and Complexity in Durham ACiD 2005*, volume 4 of *Texts in Algorithmics*, pages 1–41. King's College Publications, 2005.

[24] H. Fernau. Roman Domination: a parameterized perspective. *International Journal of Computer Mathematics*, 85:25–38, 2008.

[25] J. F. Fink and M. S. Jacobson. n-domination in graphs. In *Graph Theory and Its Applications to Algorithms and Computer Science*, pages 283–300. Wiley, 1985.

[26] F. V. Fomin, D. Lokshtanov, N. Misra, and S. Saurabh. Planar F-deletion: Approximation, kernelization and optimal FPT algorithms. In *53rd Annual IEEE Symposium on Foundations of Computer Science, FOCS*, pages 470–479. IEEE Computer Society, 2012.

[27] A. Hansberg and L. Volkmann. Upper bounds on the k-domination number and the k-Roman domination number. *Discrete Applied Mathematics*, 157:1634–1639, 2009.

[28] K. Kämmerling and L. Volkmann. Roman k-domination in graphs. *J. Korean Math. Soc.*, 46:1309–1318, 2009.

[29] P. C. B. Lam and B. Wei. On the total domination number of graphs. *Utilitas Mathematica*, 72:223–240, 2007.

[30] J. L. Mashburn, T. W. Haynes, S. M. Hedetniemi, S. T. Hedetniemi, and P. J. Slater. Differentials in graphs. *Utilitas Mathematica*, 69:43–54, 2006.

[31] B. McCuaig and B. Shepherd. Domination in graphs of minimum degree two. *Journal of Graph Theory*, 13:749–762, 1989.

[32] N. S. Narayanaswamy, V. Raman, M. S. Ramanujan, and S. Saurabh. LP can be a cure for parameterized problems. In C. Dürr and T. Wilke, editors, *29th International Symposium on Theoretical Aspects of Computer Science, STACS*, volume 14 of *LIPIcs*, pages 338–349. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2012.

[33] C. T. Nguyen, J. Shen, M. Hou, Li Sheng, W. Miller, and L. Zhang. Approximating the spanning star forest problem and its application to genomic sequence alignment. *SIAM Journal on Computing*, 38(3):946–962, 2008.

[34] P. J. Slater. Enclaveless sets and MK-systems. *Journal of Research of the National Bureau of Standards*, 82(3):197–202, 1977.